



Inferring Needless Write Memory Accesses on Ethereum Bytecode^{*}

Elvira Albert¹ , Jesús Correas¹ , Pablo Gordillo¹  ,
Guillermo Román-Díez² , and Albert Rubio¹ 

¹ Complutense University of Madrid, Madrid, Spain
pabgordi@ucm.es

² Universidad Politécnica de Madrid, Madrid, Spain

Abstract. Efficiency is a fundamental property of any type of program, but it is even more so in the context of the programs executing on the blockchain (known as *smart contracts*). This is because optimizing smart contracts has direct consequences on reducing the costs of deploying and executing the contracts, as there are fees to pay related to their bytes-size and to their resource consumption (called *gas*). Optimizing memory usage is considered a challenging problem that, among other things, requires a precise inference of the memory locations being accessed. This is also the case for the Ethereum Virtual Machine (EVM) bytecode generated by the most-widely used compiler, `solc`, whose rather unconventional and low-level memory usage challenges automated reasoning. This paper presents a static analysis, developed at the level of the EVM bytecode generated by `solc`, that infers write memory accesses that are needless and thus can be safely removed. The application of our implementation on more than 19,000 real smart contracts has detected about 6,200 needless write accesses in less than 4 hours. Interestingly, many of these writes were involved in memory usage patterns generated by `solc` that can be greatly optimized by removing entire blocks of bytecodes. To the best of our knowledge, existing optimization tools cannot infer such needless write accesses, and hence cannot detect these inefficiencies that affect both the deployment and the execution costs of Ethereum smart contracts.

1 Introduction

EVM and memory model. Ethereum [27] is considered the world-leading programmable blockchain today. It provides a virtual machine, named EVM (Ethereum Virtual Machine) [21], to execute the programs that run on the blockchain. Such programs, known as Ethereum “smart contracts”, can be written in high-level programming languages such as Solidity [6], Vyper [4], Serpent [3] or Bamboo [1] and they are then compiled to EVM bytecode. The EVM bytecode is the code finally deployed in the blockchain, and has become a uniform format to develop analysis and optimization tools. The memory model of EVM programs has been described in previous work [17, 19, 26, 27]. Mainly, there are three

^{*} This work was funded partially by the Spanish MCIU, AEI and FEDER (EU) projects PID2021-122830OB-C41 and PID2021-122830OA-C44 and by the CM project S2018/TCS-4314 co-funded by EIE Funds of the European Union.

regions in which data can be stored and accessed: (1) The EVM is a stack-based virtual machine, meaning that most instructions perform computations using the topmost elements in a machine *stack*. This memory region can only hold a limited amount of values, up to 1024 256-bit words. (2) EVM programs store data persistently using a memory region named *storage* that consists of a mapping of 256-bit addresses to 256-bit words and whose contents persist between external function calls. (3) The third memory region is a local volatile memory area that we will refer to as EVM *memory*, and which is the focus of our work. This memory area behaves as a simple word-addressed array of bytes that can be accessed by byte or as a one-word group. The EVM memory can be used to allocate dynamic local data (such as arrays or structs) and also for specific EVM bytecode instructions which have been designed to require some lengthy operands to be stored in local memory. This is the case of the instructions for computing cryptographic hashes, or for passing arguments to and returning data from external function calls. Compilers use the stack and volatile memory regions in different ways. The most-used Solidity compiler `solc` generates EVM code that uses the stack for storing value-type local variables, as well as intermediate values for complex computations and jump addresses, whereas reference-type local variables such as array types and user-defined struct types are located in memory. For instance, when a Solidity function returns a struct variable, the required memory for the struct is allocated and initialized at the beginning of the function execution. However, the allocated memory is not always accessed as we illustrate in the following function (that belongs to the contract in Fig. 1):

```

1 function _ownershipAt(uint256 i) private returns (TokenOwnership memory) {
2   return c.unpackedOwnership(_packedOwnerships[i]);
3 }

```

Although the execution of `_ownershipAt` allocates memory for the return value declared in the function definition, the execution of the function is reserving a different memory space for the actual returned struct obtained from `unpackedOwnership` and, thus, the first reservation and its initialization are needless. The focus of our work is on detecting such needless write memory accesses on the code generated by `solc`. Nevertheless, as the analysis works at EVM level, it could be easily adapted to EVM code generated by any other compiler.

Optimization. Optimization of Ethereum smart contracts is a hot research topic, see e.g. [9, 10, 12–14, 22, 24] and their references. This is because the reduction of their costs is relevant for three reasons: (1) *Deployment fees*. When the contract is deployed on the blockchain, the owner pays a fee related to the size in bytes of the bytecode. Hence, a clear optimization criterion is the bytes-size of the program. The Solidity compiler `solc` [6] has as optimization target such bytes-size reduction. (2) *Gas-metered execution*. There is a fee to be paid by each client to execute a transaction in the blockchain. This fee is a fixed amount per transaction plus the cost of executing all bytecode instructions within the function being invoked within the transaction. This cost is measured in “gas” (which is then priced in the corresponding cryptocurrency) and this is why the execution is said to be gas-metered. The EVM specification ([27] and more recent updates)

provides a precise gas consumption for each bytecode instruction in the language. The goal of most EVM bytecode optimization tools [9, 10, 12–14, 22] is to reduce such gas consumption, as this will revert on reducing the price of all transactions on the smart contract. (3) *Enlarging Ethereum’s capability*. Due to the huge volume of transactions that are being demanded, there is a huge interest in enlarging the capability of the Ethereum network to increase the number of transactions that can be handled. Optimization of EVM bytecode in general –and of its memory usage in particular– is an important step contributing into this direction.

Challenges and contributions. Optimizing memory usage is considered a challenging problem that requires a precise inference of the memory locations being accessed, and that usually varies according to the memory model of the language being analyzed, and to the compiler that generates the code to be executed. In the case of Ethereum smart contracts generated by the `solc` compiler, the memory model is rather unconventional and its low-level memory usage patterns challenge automated reasoning. On one hand, instead of having an instruction to allocate memory, the allocation is performed by a sequence of instructions that use the value stored at address `0x40` as the *free memory pointer*, i.e., a pointer to the first memory address available for allocating new memory. In the general case, the memory is structured as a sequence of *slots*: a slot is composed of several consecutive memory locations that are accessed in the bytecode from the same initial memory location plus a corresponding offset. A slot might just hold a data structure created in the smart contract but also, when nested data structures are used, from one slot we can find pointers to other memory slots for the nested components. Finally, there are other type of *transient* slots that hold temporary data and that need to be captured by a precise memory analysis as well. These features pose the main challenges to infer needless write accesses and, to handle them accurately, we make the following main contributions: (1) we present a *slot analysis* to (over-)approximate the slots created along the execution and the program points at which they are allocated; (2) we then introduce a *slot usage analysis* which infers the accesses to the different slots from the bytecode instructions; (3) we finally infer *needless write accesses*, i.e., program points where the memory is written but is never read by any subsequent instruction of the program; and (4) we implement the approach and perform a thorough experimental evaluation on real smart contracts detecting needless write accesses which belong to highly optimizable memory usage patterns generated by `solc`. Finally, it is worth mentioning that the applications of the memory analysis (points 1 and 2) go beyond the detection of needless write accesses: a precise model of the EVM memory is crucial to enhance the accuracy of any posterior analysis (see, e.g., [19] for other concrete applications of a memory analysis).

2 Memory Layout and Motivating Examples

Memory Opcodes. The EVM instruction set contains the usual instructions to access memory: the most basic instructions that operate on memory are `MLOAD`

```

4 struct TokenOwnership {
5     address addr;
6     uint64 startTs;
7     bool burned;
8 }
9
10 contract Running1 {
11     //...
12     function unpackedOwnership
13         (uint256 packed) public
14 s1 s2 returns (TokenOwnership
15             memory ownership) {
16         ownership.addr = ...;
17         ownership.startTs = ...;
18         ownership.burned = ...;
19     }
20
21 contract Running2 {
22     Running1 c;
23     mapping(uint256=>uint256) private _packedOwnerships;
24     // ...
25     function _ownershipAt(uint256 i) private
26 s3 returns (TokenOwnership memory) {
27 s4 return c.unpackedOwnership(_packedOwnerships[i]);
28 }
29
30     function explicitOwnershipOf(uint256 tokenId)
31 s5 public returns (TokenOwnership memory) {
32     TokenOwnership memory ownership;
33     if (...) { return ownership; }
34     ownership = _ownershipAt(tokenId);
35     //...
36     return ownership;
37 }
38 }

```


Fig. 1: Excerpt of smart contract ERC721A.

and `MSTORE`, which load and store a 32-byte word from memory, respectively.³ The `solc` compiler generates code to handle memory with a cumulative model in which memory is allocated along the execution of the program and is never released. In contrast to other bytecode virtual machines, like the *Java Virtual Machine*, the EVM does not have a particular instruction to allocate memory. The allocation is performed by a sequence of instructions that use the value stored at address `0x40` as the *free memory pointer*, i.e., a pointer to the first memory address available for allocating new memory. In what follows, we use $mem(x)$ to refer to the content stored in memory at location x .


Memory Slots. In the general case, memory is structured as a sequence of *slots*. A slot is composed of consecutive memory locations that are accessed by using its initial memory location, which we call the *base reference* (*baseref* for short) of the slot, plus the corresponding offset needed to access a specific location within the slot. Slots usually store (part of) some data structure created in the Solidity program (e.g., an array or a struct) and whose length can be known.

Example 1 (slots). Fig. 1 shows an excerpt of smart contract ERC721A [2] which contains two different contracts `Running1` and `Running2`. We have omitted non-relevant instructions such as those that appear at lines 15-17 (L15-L17 for short). The contract `Running1` to the left of Fig. 1 contains the public function `unpackedOwnership` that returns a struct of type `TokenOwnership` defined at L4-L7. The contract `Running2`, shown to the right, contains the public function `explicitOwnershipOf` that returns, depending on a non-relevant condition, an empty struct of type `TokenOwnership` (L29) or the `TokenOwnership` received from a call to function `unpackedOwnership` of contract `Running1` (L23), which is done in the private function `_ownershipAt`. The execution of function `unpackedOwnership` in `Running1` allocates two different memory slots at L13: s_1 , for the returned variable `ownership`, and s_2 , which is used for actually returning from the function the contents of `ownership`:

³ Although the local memory is byte addressable with instruction `MSTORE8`, to keep the description simpler, we only consider the general case of word-addressable `MSTORE`.



The function `explicitOwnershipOf` in `Running2` makes a more intensive use of the memory which can be seen in this graphical representation:



The execution of this function might create up to six different slots. At L27 and L28, it creates two slots, one for the struct declared in the `returns` part of the function header (s_3) and one for the local variable `ownership` (s_4). Depending on the evaluation of the condition in the `if` sentence, it might create the slots needed to perform the call to `_ownershipAt` and, consequently, the external call to `Running1.unpackedOwnership`. The invocation to the private function involves three slots: one for the struct declared in the `returns` part of `_ownershipAt` in L31 (s_6), one slot to manage the external call data in L23 (s_7), and one slot for storing the results of the private function `_ownershipAt` in L31 (s_8). Finally, a new slot (s_5) is created for returning the results of `explicitOwnershipOf`. This new slot might contain the contents of s_4 or s_8 , depending on the `if` evaluation.

When an amount of memory t is to be allocated, the slot reservation is made by reading and incrementing the free memory pointer (`mem(0x40)`) t positions. From this update on, the *base reference* to the slot just allocated is used, and subsequent accesses to the slot are performed by means of this baseref, possibly incremented by an offset.

Example 2 (memory slot reservation). The following excerpt of EVM code allocates a slot of type `TokenOwnership`. The EVM bytecode performs three steps:

```

(i) load the current value of the free memory pointer mem(0x40) that will be used as the
baseref of the new slot; (ii) compute the new free memory address by adding  $t$  to the baseref;
and (iii), store the new free memory pointer in mem(0x40). Additionally, in the same block
of the CFG, the slot reservation is followed by the slot initialization at 0x19A, 0x1AB and
0x1B4.
    0x175: JUMPDEST
    0x176: PUSH1 0x40
    0x178: MLOAD // (i) baseref
           DUP1
           PUSH1 0x60 // Sizeof "t"
           ADD // (ii) baseref+0x60
    0x17D: PUSH1 0x40
    0x17F: MSTORE // (iii)
           ...
    0x19A: MSTORE // baseref+0x00
           ...
    0x1AB: MSTORE // baseref+0x20
           ...
    0x1B4: MSTORE // baseref+0x40

```

Solidity reference type values such as arrays, struct typed variables and strings are stored in memory using this general pattern, with some minor differences. However, there are some cases in which the steps detailed above vary and the size of the slot is not known in advance, and thus the free memory pointer cannot be updated at this point. For instance, when data is returned by an external call, its length is unknown beforehand and hence the free memory pointer is updated only after the memory pointed to is written. In other cases, the free memory is used as a temporary region with a short lifetime, as in the case of parameter

passing to external calls, and the free memory pointer is not updated. These variants of the general schema must be detected by a precise memory analysis. To this end, we consider that a slot is in *transient* state when its baseref has been read from $mem\langle 0x40 \rangle$ but the free memory pointer has not been updated, and it is in *permanent* state when the free memory pointer has been pushed forward.

Example 3 (transient slot). Now we focus on the external call in L23 of Running2, which performs a `STATICCALL`, reading from the stack (see [27] for details) the memory location of the input arguments and the location where the results of the call will be saved. Interestingly, both locations reuse the same slot (it corresponds to s_7) as it can be seen in the following EVM bytecode from `_ownerShipAt`:

```

    PUSH4 0xb04dd20b // func. selector
    ...
    PUSH1 0x40
0x114: MLOAD // baseref transient slot
    ...
    DUP2
    MSTORE // stores func. selector
    PUSH1 0x04
    ADD // offset of funct. args.
    ... // copy func. args.
    MSTORE // stores func. args.
    ...
    PUSH 0x40
0x132: MLOAD // slot baseref
    ...
0x139: STATICCALL // external call
    ...
    PUSH1 0x40
0x151: MLOAD // slot baseref
    RETURNDATASIZE
    ...
    ADD // baseref + data size
    ...
0x15E: PUSH1 0x40
0x160: MSTORE // permanent slot

```

The call starts by reading the free memory pointer (at $0x114$) and storing at that address the arguments' data (which include the function selector as first argument). Importantly, the pointer is not pushed forward when the input arguments are written and thus the slot remains in transient state. Once the call at $0x139$ is executed, the result is written to memory from the baseref on (overwriting the locations used for the input arguments) and the slot is finally made permanent by reading the free memory pointer again ($0x151$) and updating it ($0x160$) by adding the actual return data size (`RETURNDATASIZE`).

Transient slots are also used when returning data from a public function to an external caller. In that case, the EVM code of the public function halts its execution using a `RETURN` instruction. It reads from the stack the memory location where the length and the data to be returned are located. However, it does not change $mem\langle 0x40 \rangle$ because the function code halts its execution at this point, as we can see in the EVM code of `explicitOwnershipOf` (corresponds to slot s_5):

```

    PUSH1 0x40
0x4D: MLOAD //ret slot baseref
    ...
    MSTORE // ret.addr (ret+0x00)
    ...
    MSTORE // ret.startTs (ret+0x20)
    ...
    MSTORE // ret.burned (ret+0x40)
    ...
    PUSH1 0x40
0x5A: MLOAD //ret slot revisit
    DUP1
    SWAP2 //Baseref of ret plus size
    SUB //Size of ret data
    ...
0x5E: SWAP1
0x5F: RETURN //ret returned

```

The baseref for the return slot is read (at $0x4D$) and it is used as a transient slot to write the struct contents to be returned by adding the corresponding offset for each field contained in the struct (instructions on the left column). The code on the left ends with the baseref plus the size of the stored data on top of the stack. After that, the baseref is read again (top of the right column) and the length of the returned data is computed (by subtracting the baseref to the baseref plus the size of the stored data) before calling the `RETURN` instruction.

3 Inference of Needless Write Accesses

This section presents our static inference of needless write accesses. We first provide some background in Sec. 3.1 on the type of control-flow-graph (CFG) and static analysis we rely upon. Then, the analysis is divided into three consecutive steps: (1) the slot analysis, which is introduced in Sec. 3.2, to identify the slots created along the execution and the program points at which they are allocated; (2) the slot usage analysis, presented in Sec. 3.3, which computes the read and write accesses to the different slots identified in the previous step; and (3) the detection of needless write accesses, given in Sec. 3.4, which finds those program points where there is a write access to a slot which has no read access later on.

3.1 Context-Sensitive CFG and Flow-Sensitive Static Analysis

The construction of the CFG of Ethereum smart contracts is a key part of any decompiler and static analysis tool and has been subject of previous research [15, 16, 25]. The more precise the CFG is, the more accurate our analysis results will be. In particular, context-sensitivity [16] on the CFG construction is vital to achieve precise results. Our implementation of context-sensitivity is realized by cloning the blocks which are reached from different contexts.

Example 4 (context-sensitive CFG). The EVM code of `Running2` creates multiple slots for handling structs of type `TokenOwnership`. Interestingly, all these slots are created by means of the same EVM code shown in Ex. 2, which corresponds to the CFG block that starts at program point `0x175`. As this block is reached from different contexts, the context-sensitive CFG contains three clones of this block: `0x175`, which creates s_3 at L27; `0x175_0`, which creates s_4 used at L28; and `0x175_1`, which reserves s_6 , created at L22. Block cloning means that program points are cloned as well, and we adopt the same subindex notation to refer to the program points included in the cloned block: e.g. program point `0x178` contains the `MLOAD 0x40` that gets the baseref of the slot reserved at block `0x178`, and `0x178_0` to the same `MLOAD` but at `0x178_0`, etc.

In what follows, we assume that cloning has been made and the memory analysis using the resulting CFG (with clones) is thus context-sensitive as well, without requiring additional extensions. As usual in standard analyses [23], one has to define the notion of *abstract state* which defines the abstract information gathered in the analysis and the *transfer function* which models the analysis output for each possible input. Besides context-sensitivity, the two analyses that we will present in the next two sections are *flow-sensitive*, i.e., they make a flow-sensitive traversal of the CFG of the program using as input for analyzing each block of the CFG the information inferred for its callers. When the analysis reaches a CFG block with new information, we use the operation \sqcup to join the two abstract states, and the operator \sqsubseteq to detect that a fixpoint is reached and, thus, that the analysis terminates. The operations \sqcup and \sqsubseteq , the abstract state, and transfer function, will be defined for each particular analysis.

3.2 Slot Analysis

The slot analysis aims at inferring the *abstract slots*, which are an abstraction of all memory allocations that will be made along the program execution. The

slots inferred are *abstract* because over-approximation is made at the level of the program points at which slots are allocated. Therefore, an abstract slot might represent multiple (not necessarily consecutive) real memory slots, e.g., when memory is allocated within a loop. The slot analysis will look for those program points at which the value stored in $mem\langle 0x40 \rangle$ is read for reserving memory space. These program points are relevant in the analysis for two reasons: firstly, to obtain the baseref of the memory slot, and, secondly, because from this point on, the memory reservation of the corresponding slot has started and it is pending to become permanent at some subsequent program point. The output of the slot analysis is a set which contains the allocated abstract slots, named \mathcal{S}_{all} in Def. 2 below. Each allocated abstract slot (i.e., each element in \mathcal{S}_{all}) is in turn a set of program points, as the same abstract slot might have several program points where $mem\langle 0x40 \rangle$ is read before its reservation becomes permanent. In order to obtain \mathcal{S}_{all} , the memory analysis makes a flow-sensitive traversal of the (context-sensitive) CFG of the program that keeps at every program point the set of transient slots (i.e. whose baseref has been read but it has not yet made permanent) and applies the transfer function in Def. 1 to each bytecode instruction within the blocks until a fixpoint is reached. An *abstract state* of the analysis is a set $\mathcal{S} \subseteq \wp(\mathcal{P}_R)$, where \mathcal{P}_R is the set of all program points at which $mem\langle 0x40 \rangle$ is read. The analysis of the program starts with $\mathcal{S} = \{\emptyset\}$ at all program points and takes \sqcup and \sqsubseteq as the set union and inclusion operations. Termination is trivially guaranteed as the number of program points is finite and so is $\wp(\mathcal{P}_R)$. In what follows, Ins is the set of EVM instructions and, for simplicity, we consider `MLOAD 0x40` and `MSTORE 0x40` as single instructions in Ins .

Definition 1 (slot analysis transfer function). *Given a program point pp with an instruction $I \in Ins$, an abstract state \mathcal{S} , and $\mathcal{K} = \{MSTORE\ 0x40, RETURN, REVERT, STOP, SELFDSTRUCT\}$, the slot analysis transfer function ν is defined as a mapping $\nu : Ins \times \wp(\mathcal{S}) \mapsto \wp(\mathcal{S})$ computed according to the following table:*

	I	$\nu(\mathbf{I}, \mathcal{S})$
(1)	<code>MLOAD 0x40</code>	$\{s \cup \{pp\} \mid s \in \mathcal{S}\}$
(2)	$I \in \mathcal{K}$	$\{\emptyset\}$
(3)	<i>otherwise</i>	\mathcal{S}

Let us explain intuitively how the above transfer function works. As we have seen in Sec. 2, in an EVM program all memory reservations start by reading $mem\langle 0x40 \rangle$ by means of a `MLOAD` instruction preceded by a `PUSH 0x40` instruction (case 1 in Def. 1). In this case, the transfer function adds to all sets in \mathcal{S} the current program point, since this is, in principle, an access to the same slots that were already open at this program point and are not permanent yet. To properly identify the slots, our analysis also searches for those program points at which slots reservations are made permanent (case 2 in Def. 1), i.e., those program points with instructions $I \in \mathcal{K}$. The most frequently used instruction to make a slot reservation *permanent* is a write access to $mem\langle 0x40 \rangle$ using `MSTORE`, that pushes forward the free memory pointer such that any subsequent read access to $mem\langle 0x40 \rangle$ will allocate a different slot. The rest of instructions in \mathcal{K} finalize the execution in different forms (a normal return, a forced stop, a revert execution, etc.). In all such cases, the slot needs to be considered as a permanent slot so that we can reason later on potential needless write accesses involved in it. The

set \mathcal{S} is empty after these instructions since all transient (abstract) slots are made permanent after them. We use the notation \mathcal{S}_{pp} to refer to the abstract state computed at program point pp .

Example 5 (slot analysis). The slot analysis of `Running2` starts with $\mathcal{S}_{pp}=\{\emptyset\}$ at all program points. When it reaches the block that starts at `0x175` (see Ex. 2) \mathcal{S}_{0x175} is $\{\emptyset\}$ and it remains empty until `0x178`, where the baseref of s_3 is read and hence $\mathcal{S}_{0x178}=\{\{0x178\}\}$. This slot is made permanent when the free memory pointer is updated at `0x17F`, thus having $\mathcal{S}_{0x17D}=\{\{0x178\}\}$ and $\mathcal{S}_{0x17F}=\{\emptyset\}$. Following the same pattern, s_4 and s_6 are resp. reserved at instructions `0x178_0` and `0x178_1` and closed at `0x17F_0` and `0x17F_1` (at the cloned blocks). On the other hand, the baseref of s_5 is read at two consecutive program points (`0x4D` and `0x5A`) and updated at `0x5F`, and thus, we have $\mathcal{S}_{0x4D}=\{\{0x4D\}\}$ and the same until $\mathcal{S}_{0x5A}=\{\{0x4D, 0x5A\}\}$ and again the same until $\mathcal{S}_{0x5F}=\{\emptyset\}$. Finally, after the execution of `STATICCALL` (see Ex. 3) we have three consecutive reads of `mem<0x40>` at `0x114`, `0x132` and `0x151` that refer to the same slot s_7 , which is made permanent at `0x160`. Therefore, we have $\mathcal{S}_{0x151}=\{\{0x114, 0x132, 0x151\}\}$ and $\mathcal{S}_{0x160} = \{\emptyset\}$.

Using the transfer function, as mentioned in Sec. 3.1, our analysis makes a flow-sensitive traversal of the (context-sensitive) CFG of the program that uses as input for analyzing each block the information inferred for its callers. When a fixpoint is reached, we have an abstract state for each program point that we use to compute the set of abstract slots allocated in the program, named \mathcal{S}_{all} .

Definition 2. *The set of allocated abstract slots \mathcal{S}_{all} is defined as*

$$\mathcal{S}_{all} = \bigcup_{pp \in \mathcal{P}_W} \mathcal{S}_{pp-1}, \text{ where } \mathcal{P}_W \text{ is the set of all program points } pp:I \text{ where } I \in \mathcal{K}.$$

Example 6 (\mathcal{S}_{all} computation). With the values of $\mathcal{S}_{0x17F-1}$, $\mathcal{S}_{0x17F.0-1}$, $\mathcal{S}_{0x17F.1-1}$, $\mathcal{S}_{0x160-1}$ and \mathcal{S}_{0x5F-1} from Ex. 5, at the end of the slot analysis of `Running2`, we have:

$$\mathcal{S}_{all} = \underbrace{\{\{0x178\}\}}_{s_3}, \underbrace{\{\{0x178_0\}\}}_{s_4}, \underbrace{\{\{0x178_1\}\}}_{s_6}, \underbrace{\{\{0x114, 0x132, 0x151\}\}}_{s_7}, \underbrace{\{\{0x5A, 0x4D\}, \dots\}}_{s_5}.$$

Note that, the cloning of block `0x175` allows our analysis to detect three different slots, s_3 , s_4 and s_6 , for the same program point, `0x178`, in the original EVM code.

The next example shows the behavior of the analysis when the program contains loops, and an abstraction is needed for approximating the slots.

Example 7 (loops). Fig. 2 shows the contract `Running3` that includes the function `explicitOwnershipOf` from the smart contract at [2] (made through a `STATICCALL`). This function receives an array of token identifiers as argument and returns an array of `TokenOwnership` structs that is populated invoking the function `explicitOwnershipOf` from `Running2` inside a loop. The slots identified by the analysis for contract `Running3` shown in Fig. 2 are: s_9 , which is created for making a copy of parameter `tokenIds` to memory; s_{10} , which creates the local array `ownerships` (L44) that contains the array length and pointers to the structs identified initially by s_{11} (and later on by s_{13}); s_{12} for `STATICCALL` input arguments and return data (L46); s_{13} which abstracts the structs for storing the `STATICCALL` output results (L46); and s_{14} , which includes the length of ownership

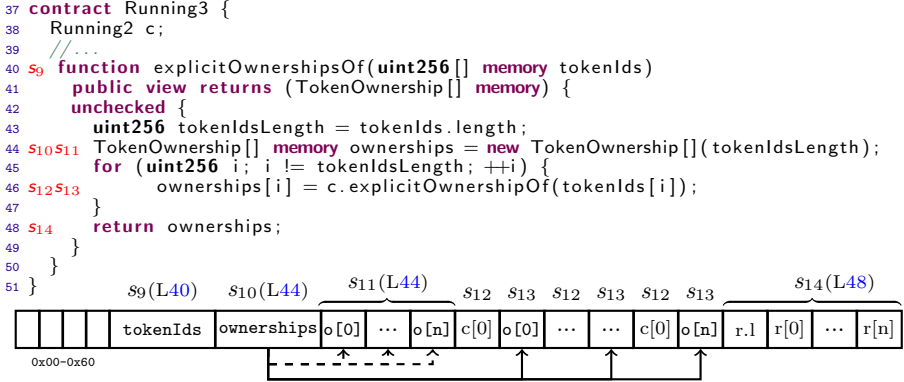


Fig. 2: Solidity code of contract `Caller`.

and a copy of s_{13} for returning the results (L48). The important point is that, the local array declaration at L44 produces a loop to allocate as many structs as elements are contained in the array. For this reason, s_{11} is an abstract slot that represents all `TokenOwnership`'s initially added to the array. Similarly, s_{12} and s_{13} are created inside the `for` loop, and each abstract slot represents as many concrete slots as iterations are performed by the loop. Note that, each iteration of the loop creates one instance of s_{12} for getting the results from the call, and it is copied later to s_{13} and pointed by `ownerships` (s_{10}).

As notation, we will use a unique numeric identifier (1, 2, ...) to refer to each abstract slot (represented in \mathcal{S}_{all} as a set) and retrieve it by means of function $get_id(a), a \in \mathcal{S}_{all}$. We use \mathcal{A} to refer to the set of all such identifiers in the program. Also, given a program point pp with an instruction `MLOAD 0x40`, we define the function $get_slots(pp)$ to retrieve the identifiers of the elements of \mathcal{S}_{all} that might be referenced at pp as follows: $get_slots(pp) = \{id \mid a \in \mathcal{S}_{all} \wedge pp \in a \wedge id = get_id(a)\}$.

3.3 Slot Access Analysis

While Sec. 3.2 looked for allocations, the next step of the analysis is the inference of the program points at which the inferred abstract slots might be accessed. To do so, our slot access analysis needs to propagate the references to the abstract slots that are saved at the different positions of the execution stack. Importantly, we keep track, not only of the stack positions, but also, in order to abstract complex data structures stored in memory (e.g., arrays of structs), we need to keep track of the abstract slots that could be saved at memory locations. As seen in Ex. 7, a memory location within a slot might contain a pointer to another memory location of another slot, as it happens when nested data structures are used. Thus, an abstract state is a mapping at which we store the potential slots saved at stack positions or at memory locations within other slots.

Definition 3 (memory analysis abstract state). A memory analysis abstract state is a mapping π of the form $\mathcal{T} \cup \mathcal{A} \mapsto \wp(\mathcal{A})$.

\mathcal{T} is the set containing all stack positions, which we represent by natural numbers from 0 (bottom of the stack) on, and \mathcal{A} is the set of abstract slots identifiers computed in Sec. 3.2. We refer to the set of all memory analysis abstract states as AS . Note that, for each entry, we keep a set of potential slots for each stack position because a block might be reached from several blocks with different execution stacks, e.g., in loops or *if-then-else* structures. In what follows, we assume that, given a value k , the map π returns the empty set when $k \notin \text{dom}(\pi)$. The inference is performed by a flow-sensitive analysis (as described in Sec. 3.1) that keeps track of the information about the abstract slots used at any program point by means of the following transfer function.

Definition 4 (memory analysis transfer function). *Given an instruction I with n input operands at program point pp and an abstract state π , the memory analysis transfer function τ is defined as a mapping $\tau: \text{Ins} \times AS \mapsto AS$ of the form:*

	I	$\tau(I, \pi)$		I	$\tau(I, \pi)$
(1)	<code>MLOAD 0x40</code>	$\pi[t \mapsto \text{get_slots}(pp)]$	(4)	<code>SWAPi</code>	$\pi[t \mapsto \pi(t-i), t-i \mapsto \pi(t)]$
(2)	<code>MLOAD</code>	$\pi[t \mapsto \{m \mid s \in \pi(t) \wedge m \in \pi(s)\}]$	(5)	<code>DUPi</code>	$\pi[t+1 \mapsto \pi(t-i+1)]$
(3)	<code>MSTORE</code>	$\pi[s \mapsto \pi(s) \cup \pi(t-1)] \setminus \{t, t-1\} \forall s \in \pi(t)$	(6)	<i>otherwise</i>	$\pi \setminus x \quad t-n < x \leq t$

$t = \text{top}(pp)$ is the numerical position of the top of the stack before executing I .

Let us explain the above definition. The transfer function distinguishes between two different types of `MLOAD`: (1) accesses to location `mem(0x40)`, which return the baseref of the slots that might be used, taking them from the previous analysis through `get_slots(p)`; and (2) other `MLOAD` instructions, which could potentially return slot baserefs from memory locations. Therefore, we have to consider two possibilities: if we are reading a memory location which reads a generic value (e.g. a number) then $\pi(t) = \emptyset$; if we are reading a memory location that might store an abstract slot, then $\pi(t)$ contains all abstract slots that might be stored at that memory location. Regarding (3), `MSTORE` has two operands: the operand at t is the memory address that will be modified by `MSTORE`, and the operand at $t-1$ is the value to be stored in that address. For each element s in $\pi(t)$, the analysis adds the abstract slots that are in $\pi(t-1)$. Other instructions that are also treated by the analysis are `SWAP*` and `DUP*` shown in (4-5), that exchange or copy the elements of the stack that take part in the operation. Finally, all other operations delete the elements of the stack that are no longer used based on the number of elements taken and written to the stack (case 6).

Example 8 (transfer). Now we focus on the analysis of block `0x175`, shown in Fig. 3. As we have already explained, this block is responsible for creating the memory needed to work with several structs of type `TokenOwnership` and it is thus cloned in the CFG. In particular, we focus on the clone `0x175_1`. The analysis of the block starts with a stack of size 7 and includes at positions 3 and 4, the abstract slots s_3 and s_4 , which were created at `L26` and `L27` of Fig. 1. At `0x178_1`, `mem(0x40)` is read, and, by means of `get_slots(0x178.1)` and, considering that $\text{top}(0x178.1) = 8$, we add to π a new entry $8 \mapsto s_6$. At `0x179_1`, `0x180_1`, `0x1AA_1`, `0x1B3_1` the transfer function duplicates a slot identifier stored in the stack. `MSTORE` and `POP` instructions of the example remove a slot identifier from the stack.

PP	Instr	π	PP	Instr	π
Ox175.1	JUMPDEST	$\{3 \mapsto s_3, 4 \mapsto s_4\}$	Ox19A.1	MSTORE	$\{3 \mapsto s_3, 4 \mapsto s_4, 8 \mapsto s_6, 9 \mapsto s_6\}$
Ox176.1	PUSH1 0x40	$\{3 \mapsto s_3, 4 \mapsto s_4\}$...		
Ox178.1	MLOAD	$\{3 \mapsto s_3, 4 \mapsto s_4, 8 \mapsto s_6\}$	Ox1A9.1	AND	$\{3 \mapsto s_3, 4 \mapsto s_4, 8 \mapsto s_6, 9 \mapsto s_6\}$
Ox179.1	DUP1	$\{3 \mapsto s_3, 4 \mapsto s_4, 8 \mapsto s_6, 9 \mapsto s_6\}$	Ox1AA.1	DUP2	$\{3 \mapsto s_3, 4 \mapsto s_4, 8 \mapsto s_6, 9 \mapsto s_6, 11 \mapsto s_6\}$
Ox17A.1	PUSH1 0x60	$\{3 \mapsto s_3, 4 \mapsto s_4, 8 \mapsto s_6, 9 \mapsto s_6\}$	Ox1AB.1	MSTORE	$\{3 \mapsto s_3, 4 \mapsto s_4, 8 \mapsto s_6, 9 \mapsto s_6\}$
Ox17C.1	ADD	$\{3 \mapsto s_3, 4 \mapsto s_4, 8 \mapsto s_6, 9 \mapsto s_6\}$...		
Ox17D.1	PUSH1 0x40	$\{3 \mapsto s_3, 4 \mapsto s_4, 8 \mapsto s_6, 9 \mapsto s_6\}$	Ox1B2.1	ISZERO	$\{3 \mapsto s_3, 4 \mapsto s_4, 8 \mapsto s_6, 9 \mapsto s_6\}$
Ox17F.1	MSTORE	$\{3 \mapsto s_3, 4 \mapsto s_4, 8 \mapsto s_6\}$	Ox1B3.1	DUP2	$\{3 \mapsto s_3, 4 \mapsto s_4, 8 \mapsto s_6, 9 \mapsto s_6, 11 \mapsto s_6\}$
Ox180.1	DUP1	$\{3 \mapsto s_3, 4 \mapsto s_4, 8 \mapsto s_6, 9 \mapsto s_6\}$	Ox1B4.1	MSTORE	$\{3 \mapsto s_3, 4 \mapsto s_4, 8 \mapsto s_6, 9 \mapsto s_6\}$
...			Ox1B5.1	POP	$\{3 \mapsto s_3, 4 \mapsto s_4, 8 \mapsto s_6\}$
Ox198.1	AND	$\{3 \mapsto s_3, 4 \mapsto s_4, 8 \mapsto s_6, 9 \mapsto s_6\}$	Ox1B6.1	SWAP1	$\{3 \mapsto s_3, 4 \mapsto s_4, 7 \mapsto s_6\}$
Ox199.1	DUP2	$\{3 \mapsto s_3, 4 \mapsto s_4, 8 \mapsto s_6, 9 \mapsto s_6, 11 \mapsto s_6\}$	Ox1B7.1	JUMP	$\{3 \mapsto s_3, 4 \mapsto s_4, 7 \mapsto s_6\}$

Fig. 3: Block of the CFG that reserves memory slot for struct

As it is flow-sensitive, the analysis of each block of the CFG takes as input the join \sqcup of the abstract states computed with the transfer function for the blocks that jump to it, and keeps applying the memory analysis transfer function until a fixpoint is reached. The operation $A \sqcup B$ is the result of joining, by means of operation \cup , all entries from maps A and B . Operation \sqsubseteq is defined as expected, $A \sqsubseteq B$, when B includes entries that are not in $dom(A)$ or when we have an entry $v \in dom(A) \cap dom(B)$ such that $A(v) \subseteq B(v)$. Again, termination of the computation is guaranteed because the domain is finite.

Example 9 (joining abstract states). The EVM code of `explicitOwnershipOf` of Fig. 1 uses s_5 in both `return` sentences at L29 and L33 (see Ex. 1). This EVM code has a single return block which is reachable from two different paths from the `if` statement, and which come with different abstract states: (1) the path that corresponds to L29 comes with $\pi = \{3 \mapsto s_8\}$, and the other path (L33) with $\pi = \{3 \mapsto s_4\}$. Our analysis joins both abstract states resulting in $\pi = \{3 \mapsto \{s_4, s_8\}\}$. Because of this join, we get that the `RETURN` instruction that comes from lines L29 and L33 might return the content of the slots s_4 or s_8 .

When the fixpoint is reached, the analysis has computed an abstract state for each program point pp , denoted by π_{pp} in what follows.

Example 10 (complex data structures). The analysis of the code at Fig. 2 shows how it deals with data structures that might contain pointers to other structures, e.g. `ownerships`. The abstract slot that represents variable `ownerships` is s_{10} , which is written, by means of `MSTORE` at two program points, say pp_1 and pp_2 which, resp., come from L44 and L46 of the Solidity code. The input abstract state that reaches pp_1 is $\{2 \mapsto s_9, 6 \mapsto s_{10}, 8 \mapsto s_{10}, 9 \mapsto s_{11}, 10 \mapsto s_{10}\}$, and the transfer function of `MSTORE` leaves the abstract state as $\pi_{pp_1} = \{2 \mapsto s_9, 6 \mapsto s_{10}, 8 \mapsto s_{10}, s_{10} \mapsto s_{11}\}$. At this point, we can see that variable `ownerships` is initialized with empty structs and, to represent it, our analysis includes in π the entry $s_{10} \mapsto s_{11}$ as it is described in instruction `MSTORE` of the transfer function at Def. 4. The second write to s_{10} is performed by another `MSTORE` instruction at pp_2 . The input abstract state for pp_2 is $\{2 \mapsto s_9, 5 \mapsto s_{10}, 7 \mapsto s_{13}, 8 \mapsto s_{13}, 9 \mapsto s_{10}, s_{10} \mapsto s_{11}\}$, and thus we get $\pi_{pp_2} = \{2 \mapsto s_9, 5 \mapsto s_{10}, 7 \mapsto s_{13}, s_{10} \mapsto \{s_{11}, s_{13}\}\}$. Interestingly, at pp_2 , we detect that s_{11} might also store the structs returned by the call to `c.explicitOwnershipOf(tokenIds[i])`, identified by s_{13} , which is added to

$s_{10} \mapsto \{s_{11}, s_{13}\}$. Finally, s_{10} is read at the end of the method, returning the set $\{s_{11}, s_{13}\}$, to copy the content of `ownerships` to s_{14} , the slot used in the return.

3.4 Inference of Needless Write Memory Accesses

With the results of the previous analysis, we can compute the maps \mathcal{R} and \mathcal{W} , which are of the form $pp \mapsto \wp(\mathcal{A})$ and capture the slots that might be read or written, resp., at the different program points. To do so, as multiple EVM instructions, e.g. `RETURN`, `CALL`, `LOG`, `CREATE`, ..., might read, or write, memory locations taking the concrete location from the stack, we define functions $mr(I)$ and $mw(I)$ that, given an EVM instruction I , return the position in the stack of the address to be read and written by I , resp. If the instruction does not read/write any memory position, function $mr(I) = \perp / mw(I) = \perp$. For example, $mr(\text{MLOAD}) = 0$ as it reads the top of the stack and $mw(\text{MLOAD}) = \perp$, or $mr(\text{STATICCALL}) = 2$ and $mw(\text{STATICCALL}) = 4$. Now, we define the read/write maps \mathcal{R}/\mathcal{W} :

Definition 5 (memory read/write accesses map). *Given an EVM program P , such that $pp \equiv I \in P$ and being $t = \text{top}(pp)$, we define maps \mathcal{R} and \mathcal{W} as follows:*

$$\mathcal{R}(pp) = \begin{cases} \emptyset & mr(I) = \perp \\ \pi_{pp-1}(t - mr(I)) & \text{otherwise} \end{cases} \quad \mathcal{W}(pp) = \begin{cases} \emptyset & mw(I) = \perp \\ \pi_{pp-1}(t - mw(I)) & \text{otherwise} \end{cases}$$

Example 11 (\mathcal{R}/\mathcal{W} maps). Let us illustrate the computation of $\mathcal{R}(0x139)$ and $\mathcal{W}(0x139)$, which contains the `STATICCALL` of `Running2`. With the analysis information obtained from the analysis we have that $\text{top}(0x139) = 16$ and $\pi_{0x138} = \{3 \mapsto s_3, 4 \mapsto s_4, 7 \mapsto s_6, 10 \mapsto s_7, 12 \mapsto s_7, 14 \mapsto s_7\}$, thus we get $\mathcal{R}(0x139) = \{s_7\}$ and $\mathcal{W}(0x139) = \{s_7\}$, i.e., the slot used for managing the input and the output of the external call. Analogously, we get that $\mathcal{R}(0x178) = \{s_3\}$ and $\mathcal{W}(0x178) = \emptyset$.

The last step of our analysis consists in searching for write accesses to slots which will never be read later. To do so, we use the information computed in \mathcal{R} and \mathcal{W} . Given the CFG of the program and two program points p and $p2$, we define function $\text{reachable}(p, p2)$, which returns *true* when there exists a path in the CFG from p to $p2$. We define the set *write leaks* \mathcal{N} as follows:

Definition 6. *Given an EVM program and its \mathcal{W} and \mathcal{R} , we define \mathcal{N} as*

$$\mathcal{N} = \{pw:s \mid pw \in P \wedge s \in \mathcal{W}(pw) \wedge \neg \text{exists_read}(pw, s)\}$$

where $\text{exists_read}(pw, s) \equiv \exists pr \in \text{dom}(\mathcal{R}) \mid s \in \mathcal{R}(pr) \wedge \text{reachable}(pw, pr)$.

Intuitively, the set \mathcal{N} contains those write accesses, taken from \mathcal{W} , that are never read by subsequent blocks in the CFG. As both function reachable and the sets \mathcal{W} and \mathcal{R} are over-approximations, the computation of \mathcal{N} provides us those write accesses that can be safely removed, as the next example shows.

Example 12. Our analysis detects that at program points `0x19A`, `0x1AB` and `0x1B4` there are `MSTORE` operations that are never read in the subsequent blocks of the CFG. Such operations correspond to the memory initialization of s_3 , which is performed at `L27` of the code of Fig. 1 (see Ex. 2). Given that these write accesses are the only use of the slot, the whole reservation can be safely removed.

Moreover, the analysis detects that program points `0x19A_1`, `0x1AB_1` and `0x1B4_1`, which correspond to the reservation of s_6 performed at L22, are detected as needless. In essence, it means that s_3 and s_6 are allocated and initialized but are never used in the program. Note that, all these program points belong to two blocks cloned: (`0x175` and `0x175_1`). However, the three `MSTORE` operations of the other clone of the same block (`0x175_0`), which correspond to the allocation at L28 are not identified as non-read, as they might be used in the return of the function. For this, the precision of the context-sensitive CFG is necessary to identify these `MSTORE` operations as needless. As a result we cannot eliminate the block because it is needed in one of the clones, but still we can achieve an important optimization on the EVM code by removing the unconditional jumps to this block in the other two cases that would avoid completely the execution of all these instructions (and their corresponding gas consumption [27]).

The soundness of slots and slots access analyses states that, for each concrete slot, there exists an abstract slot in \mathcal{S}_{all} that represents it and, that any access to memory is approximated by an inferred abstract slot. Technical details can be found in an extended report [8].

4 Experimental Evaluation

This section reports on the results of the experimental evaluation of our approach, as described in Sec. 3. All components of the analysis are implemented in Python, are open-source, and can be downloaded from `github` where detailed instructions for its installation and usage are provided⁴. We use external components to build the CFGs (as this is not a contribution of our work). Our analysis tool accepts smart contracts written in versions of Solidity up to 0.8.17 and bytecode for the Ethereum Virtual Machine v1.10.25⁵. The experiments have been performed on an AMD Ryzen Threadripper PRO 3995WX 64-cores and 512 GB of memory, running Debian 5.10.70. In order to experimentally evaluate the analysis, we pulled from `etherscan.io` [5] the Ethereum contracts bound to the last 5,000 open-source verified addresses whose source code was available on July 14, 2022. From those addresses, the code of 2.18% of them raises a compilation error from `solc`. For the code bound to the 4,891 remaining addresses, the generation of the CFG (which is not a contribution of this work) timeouts after 120s on 626 of them. Removing such failing cases, we have finally analyzed 19,199 smart contracts, as each address and each Solidity file may contain several contracts in it. Note that 84.86% of the contracts are compiled with the `solc` version 0.8, presumably with the most advanced compilation techniques. The whole dataset used will be found at the above `github` link.

In order to be in a worst-case scenario for us, we run the memory analysis after executing the `solc` optimizer, i.e, we analyze bytecode whose memory usage may have been optimized already by the optimizer available in `solc`. This will allow us also to see if we can achieve further optimization with our

⁴ https://github.com/costa-group/EthIR/tree/memory_optimizer/ethir

⁵ The latest versions released up to Oct 2022.

approach. Unfortunately, we have not been able to apply our tool after running the super-optimizer GASOL [9], because it does not generate the optimized bytecode but rather it only reports on the gas and/or size gains for each of the blocks. Nevertheless, a detailed comparison of the techniques that GASOL applies and ours is given in Sec. 5, where we justify that GASOL will not find any of our needless accesses. From the 19,199 analyzed contracts, the analysis infers 679,517 abstract memory slots and detects 6,242 needless write memory accesses in 12,803s. These needless accesses occur within the code bound to 780 different addresses, i.e., 15.95% of the analyzed ones.

We have computed the number of needless accesses identified by our analysis grouped by function and the number of different contracts that contain these functions. Some of them such as `transferFrom` (1736 accesses in 439 contracts), `transfer` (1745 accesses in 441 contracts), `reflectionFromToken` (105 accesses in 6 contracts) or `withdraw` (54 accesses in 32 contracts) are functions widely used in the implementation of contracts based on ERC tokens. A manual inspection of the 10 most common public functions with the needless accesses inferred has revealed two different sources for them: some of the needless accesses are due to inefficient programming practices, while others are generated by the compiler and could be improved. As regards compiler inefficiencies, we detected bytecode that allocates memory slots that are inaccessible and cannot be used because the baseref to access them is not maintained in the stack. For example, when a struct is returned by a function, it always allocates memory for this data. However, if the return variable is not named in the header of the function, the compiler allocates memory for this data although it will never be accessed. If programmers are aware of this behavior they can avoid such generation of useless memory but, even better, this memory usage patterns can be changed in the compiler. For instance, it is reflected in L22 and L27 in Fig. 1, where the functions do not name the return variable. Hence, the compiler allocates memory for these *anonymous* data structures which are never used. Similarly, there are various situations involving external calls in which the compiler creates memory that is never used. When there is an external call that does not retrieve any result, the compiler creates two memory slots, one for retrieving the result from the call, and another one for copying a potential result to a memory variable that is never used. Finally, the compiler also creates memory that is never used for low-level plain calls for currency transfer. Even though the contract code does not use the second result returned by the low-level call, the compiler generates code for retrieving it. All these potential optimizations have been detected by means of our inference of needless write accesses and will be communicated to the `solc` developers.

5 Conclusions and Related Work

We have proposed a novel memory analysis for Ethereum smart contracts and have applied it to infer needless write memory accesses. The application of our implementation over more than 19,000 real smart contracts has detected some compilation patterns that introduce needless write accesses and that can be easily

changed in the compiler to generate more efficient code. Let us discuss related work along two directions: (1) memory analysis and (2) memory optimization. Regarding (1), we can find advanced points-to analysis developed for Java-like languages [7, 11, 18, 20]. Focusing on EVM, the static modeling of the EVM memory in [16] has some similarities with the memory analysis presented in Secs. 3.2 and 3.3, since in both cases we are seeking to model the memory although with different applications in mind. There are differences on one hand on the type of static analysis used in both cases: [16] is based on a Datalog analysis while we have defined a standard transfer function which is used within a flow-sensitive analysis. More importantly, there are differences on the precision of both analyses. We can accurately model the memory allocated by nested data structures in which the memory contains pointers to other memory slots, while [16] does not capture such type of accesses. This is fundamental to perform memory optimization since, as shown in the running examples of the paper, it allows detecting needless write accesses that otherwise would be missed. Finally, the application of the memory analysis to optimization is not studied in [16], while it is the main focus of our work.

As regards (2), optimizing memory usage is a challenging research problem that requires to precisely infer the memory positions that are being accessed. Such positions sometimes are statically known (e.g., when accessing the EVM free memory pointer) but, as we have seen, often a precise and complex inference is required to figure out the slot being accessed at each memory access bytecode. Recent work within the super-optimizer GASOL [9] is able to perform some memory optimizations at the level of each block of the CFG (i.e., intra-block). There are three fundamental differences between our work and GASOL: First, GASOL can only apply the optimizations when the memory locations being addressed refer to the same constant direction. In other words, there is no real memory analysis (namely Secs. 3.2 and 3.3). Second, the optimizations are applied only at an intra-block level and hence many optimization opportunities are missed. These two points make a fundamental difference with our approach, since detected optimizable patterns (see Sec. 4) require inter-block analysis and a precise slot access analysis, and hence cannot be detected by GASOL.

Finally, as mentioned in Sec. 1, in addition to dynamic memory, smart contracts also use a persistent memory called storage. Regarding the application of our approach to infer needless accesses in storage, there are two main points. First, there is no need to develop a static analysis to detect the slots in storage, as they are statically known (hence our inference in Sec. 3.2 and 3.3 is not needed), i.e., one can easily know the read and write sets of Def. 6. Thus, the read and write sets of our analysis can be easily defined for storage. The second point is that, as storage is persistent memory, a write storage access is not removable even if there is no further read access within the smart contract, as it needs to be stored for a future transaction. The removable write storage accesses are only those that are rewritten and not read in-between the two write accesses. Including this in our implementation is straightforward. However, this situation is rather unusual, and we believe that very few cases would be found and hence little optimization can be achieved.

References

1. Bamboo. <https://github.com/pirapira/bamboo>.
2. ERC721A. <https://etherscan.io/address/0xfcd5c0ef90715dc052dad6de08efda758aa09f60#code>.
3. Serpent. <https://github.com/ethereum/wiki/wiki/Serpent>.
4. Vyper. <https://github.com/ethereum/vyper>.
5. Etherscan. <https://etherscan.io>, 2018.
6. Solidity documentation, 2021. <https://docs.soliditylang.org/en/latest/index.html>.
7. Elvira Albert, Puri Arenas, Jesús Correas, Samir Genaim, Miguel Gómez-Zamalloa, and Germán Puebla and Guillermo Román-Díez. Object-Sensitive Cost Analysis for Concurrent Objects. *Software Testing, Verification and Reliability*, 25(3):218–271, 2015.
8. Elvira Albert, Jesús Correas, Pablo Gordillo, Guillermo Román-Díez, and Albert Rubio. Inferring needless write memory accesses on ethereum bytecode (extended version), 2023. arXiv:2301.04757 [cs.PL].
9. Elvira Albert, Pablo Gordillo, Alejandro Hernández-Cerezo, and Albert Rubio. A Max-SMT Superoptimizer for EVM handling Memory and Storage. In Dana Fisman and Grigore Rosu, editors, *28th International Conference on Tools and Algorithms for the Construction and Analysis of Systems, TACAS 2022. Proceedings*, volume 13243 of *Lecture Notes in Computer Science*, pages 201–219. Springer, 2022.
10. T. Brandstätter, S. Schulte, J. Cito, and M. Borkowski. Characterizing Efficiency Optimizations in Solidity Smart Contracts. In *2020 IEEE International Conference on Blockchain (Blockchain)*, pages 281–290, 2020.
11. Ramkrishna Chatterjee, Barbara G. Ryder, and William Landi. Relevant context inference. In Andrew W. Appel and Alex Aiken, editors, *POPL '99, Proceedings of the 26th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, San Antonio, TX, USA, January 20-22, 1999*, pages 133–146. ACM, 1999.
12. Ting Chen, Youzheng Feng, Zihao Li, Hao Zhou, Xiapu Luo, Xiaoqi Li, Xiuzhuo Xiao, Jiachi Chen, and Xiaosong Zhang. Gaschecker: Scalable analysis for discovering gas-inefficient smart contracts. *IEEE Transactions on Emerging Topics in Computing*, PP(99):1–14, 03 2020.
13. Ting Chen, Zihao Li, Hao Zhou, Jiachi Chen, Xiapu Luo, Xiaoqi Li, and Xiaosong Zhang. Towards saving money in using smart contracts. In *Proceedings of the 40th International Conference on Software Engineering: New Ideas and Emerging Results, ICSE (NIER) 2018, Gothenburg, Sweden, May 27 - June 03, 2018*, pages 81–84, 2018.
14. Bo Gao, Siyuan Shen, Ling Shi, Jiaying Li, Jun Sun, and Lei Bu. Verification assisted gas reduction for smart contracts. In *28th Asia-Pacific Software Engineering Conference, APSEC 2021, Taipei, Taiwan, December 6-9, 2021*, pages 264–274. IEEE, 2021.
15. Neville Grech, Lexi Brent, Bernhard Scholz, and Yannis Smaragdakis. Gigahorse: thorough, declarative decompilation of smart contracts. In Joanne M. Atlee, Tevfik Bultan, and Jon Whittle, editors, *Proceedings of the 41st International Conference on Software Engineering, ICSE 2019, Montreal, QC, Canada, May 25-31, 2019*, pages 1176–1186. IEEE / ACM, 2019.
16. Neville Grech, Sifis Lagouvardos, Ilias Tsatiris, and Yannis Smaragdakis. Elipmoc: Advanced decompilation of ethereum smart contracts. *Proc. ACM Program. Lang.*, 6(OOPSLA):77:1–77:27, 2022.

17. Ákos Hajdu and Dejan Jovanovic. Smt-friendly formalization of the solidity memory model. In Peter Müller, editor, *Programming Languages and Systems - 29th European Symposium on Programming, ESOP 2020, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2020, Dublin, Ireland, April 25-30, 2020, Proceedings*, volume 12075 of *Lecture Notes in Computer Science*, pages 224–250. Springer, 2020.
18. George Kastrinis and Yannis Smaragdakis. Hybrid context-sensitivity for points-to analysis. In Hans-Juergen Boehm and Cormac Flanagan, editors, *ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '13, Seattle, WA, USA, June 16-19, 2013*, pages 423–434. ACM, 2013.
19. Sifis Lagouvardos, Neville Grech, Ilias Tsatiris, and Yannis Smaragdakis. Precise static modeling of ethereum "memory". *Proc. ACM Program. Lang.*, 4(OOPSLA):190:1–190:26, 2020.
20. Ana Milanova, Atanas Rountev, and Barbara G. Ryder. Parameterized Object Sensitivity for Points-to Analysis for Java. *ACM Transactions on Software Engineering Methodology*, 14:1–41, 2005.
21. Mayukh Mukhopadhyay. *Ethereum Smart Contract Development*. Packt publishing, 2018.
22. Julian Nagele and Maria A Schett. Blockchain superoptimizer. In *Preproceedings of 29th International Symposium on Logic-based Program Synthesis and Transformation (LOPSTR 2019)*, 2019.
23. F. Nielson, H. R. Nielson, and C. Hankin. *Principles of Program Analysis*. Springer, 1999.
24. Maria A. Schett and Julian Nagele. Populating the Peephole Optimizer of a Smart Contract Compiler. In Bruno Bernardo and Diego Marmosoler, editors, *2nd Workshop on Formal Methods for Blockchains (FMBC 2020)*, volume 84 of *OpenAccess Series in Informatics (OASICs)*, pages 3:1–3:15. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2020.
25. Clara Schneidewind, Ilya Grishchenko, Markus Scherer, and Matteo Maffei. eThor: Practical and provably sound static analysis of ethereum smart contracts. In Jay Ligatti, Xinming Ou, Jonathan Katz, and Giovanni Vigna, editors, *CCS '20: 2020 ACM SIGSAC Conference on Computer and Communications Security, USA, November 9-13, 2020*, pages 621–640. ACM, 2020.
26. Petar Tsankov, Andrei Marian Dan, Dana Drachler-Cohen, Arthur Gervais, Florian Bünzli, and Martin T. Vechev. Securify: Practical Security Analysis of Smart Contracts. In David Lie, Mohammad Mannan, Michael Backes, and XiaoFeng Wang, editors, *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security, CCS 2018, Toronto, ON, Canada, October 15-19, 2018*, pages 67–82. ACM, 2018.
27. Gavin Wood. *Ethereum: A secure decentralised generalised transaction ledger*, 2019.

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

